Intelligent Robots and Computer Vision XIV Algorithms, Techniques, Active Vision, and Materials Handling,
Proc. SPIE 2588, Philadelphia, PA, Oct 23-26, 1995.

Mars lander robotics and machine vision capabilities for *in situ* planetary science

P. S. Schenker, D. L. Blaney, D. K Brown, Y. Bar-Cohen, S-S Lih, R. A. Lindemann, E. D. Paljug, J. T. Slostad, G. K. Tharp, C. F. Tucker, C. J. Voorhees, and C. Weisbin, Jet Propulsion Lab.;
E. T. Baumgartner, Mich. Tech Univ.; J. B. Singer, R. Reid, Univ. of Arizona

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive/ MS 198-219
Pasadena, CA 91109
Email: schenker@telerobotics.jpl.nasa.gov

Mechanical Engineering Department
Michigan Technological University
1400 Townsend Drive
Houghton, MI 49931

University of Arizona
Space Sciences Bldg. #92
University of Arizona
Tucson, AZ 85721

ABSTRACT

We overview our recent progress in lander-based robotics for Mars planetary science. Utilizing a 1:1 scale laboratory replica of the NASA Mars Surveyor '98 mission, JPL engineers and Mars science colleagues have demonstrated approaches to lander science functions such as robotic sample acquisition and deposition, end effector based microscopic viewing, hand-carried science instrument data collection, and science instrument emplacement by a robot. Some of the significant technical advances underlying this simulated Mars flight capability include JPL's innovation of a new lightweight, mechanically stiff, gas deployed telescopic two meter robot arm, and cooperative engineering work with Michigan Tech colleagues on automated visual positioning control of robotic sampling. University of Arizona and JPL have further developed complementary advances in lander-based imaging spectroscopy and its robotic enablement. We outline this work, summarizing its key technical features, and illustrating experimental progress with photographs and an accompanying conference videotape.

keywords: robotics, manipulators, machine vision, visual calibration, visual servoing, Mars science, Mars Surveyor Program

1. INTRODUCTION

We have begun development of robotic techniques for lander-based Mars surface science. Following earlier NASA Viking I and II missions [1], such scientific explorations will reveal fundamental information on Mars geology, climatology, and mineralogy, giving insight to the planet's history, formation, and relationship to our own [2]. The focus of our current robotics work is to expand the scope and improve the efficiency of surface science operations that can be carried out from a Mars lander. Consistent with goals of future NASA missions, the mass, volume, and power constraints for such a landed science payloads are severe -- e.g., respectively 20 kg, 70 liters, and 25 watts for a first Mars Surveyor '98 mission. Consequently, per **Figure 1**, the operational challenges to robotics are exceptional: a nominal two meters reach into the lander near field, less than 5 kg flight mass and 10 watts average power utilization. Robot volume and mass must accordingly scale to <5% that of mid-1970's Viking era technology [3]. Further, such robotic operations must be implemented in the context of long-duration autonomy under high-level time delay operator sequencing, in temperatures as low as -50 to -100 degrees C (polar seasonal extremes).

We overview our recent progress in lander-based robotics for Mars planetary science. Utilizing a 1:1 scale laboratory replica of the NASA Mars Surveyor '98 mission, JPL engineers and Mars science colleagues have demonstrated approaches to lander science functions such as robotic sample acquisition and deposition, end-effector based microscopic viewing, hand-carried science instrument data collection, and science instrument emplacement by a robot. Some of the significant technical advances underlying this simulated Mars flight capability include JPL's introduction of a new lightweight, mechanically stiff, gas deployed telescopic two meter robot arm, and cooperative engineering work with Michigan Tech colleagues on automated visual positioning control of robotic sampling. University of Arizona and JPL have jointly developed complementary advances in lander-based imaging spectroscopy and its robotic enablement. We outline this work summarizing its key technical features, and illustrating experimental progress with photographs and an accompanying conference videotape.

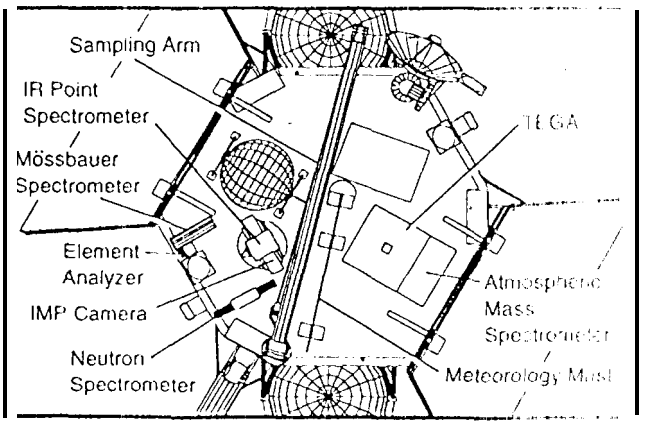

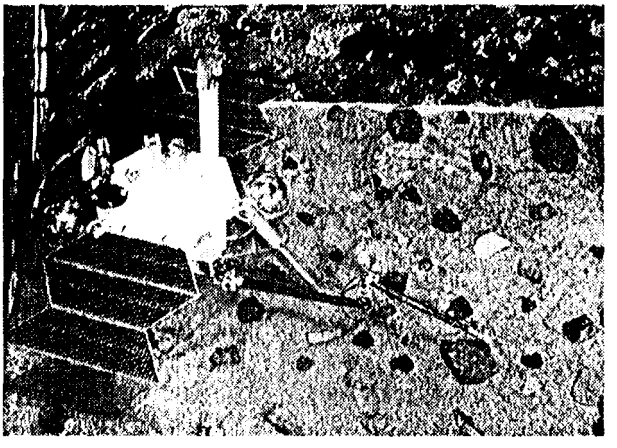

'98 Science Payload

MASS:	*20 kg*
VOLUME:	*70 liters*
POWER:	*25 W day (avg)* **10 W night (avg)**
THERMAL:	**lander provides**
COMPUTING:	**lander provides**
COM/DATA:	**RS-232, MIL-STD 1553,..**
LIFETIME:	**> 1 Mars yr (low-lat.)** **1/3 Mars yr. (high-lat.)**
DATA **VOL.:**	**10Mb / sol**
G-LOADING:	**< 50 g's all phases**

Figure 1 *Mars Lander Based Science Operations and JPL Robotic Simulations*
(top left) *a deck configuration for Mars lander science illustrating instrumentation options [4]*
(upper left) *a summary list of Mars Surveyor '98 landed mission operating constraints [4]*
(upper right) *the JPL replica Mars'98 lab simulation with 5-dof composite arm*
(bottom right) *simulated operations of lander (G. Tharp)*

2. A NEW ROBOT FOR MARS LANDER SCIENCE OPERATIONS

As noted above., a primary concern in robotics development for future NASA missions is reduction of mass and volume. These constraints arise from logistics of maximized landed science mass with minimum launch vehicle requirements, thereby enabling more frequent missions at reduced cost. To this end, we have sought ways to achieve Viking era like arm scale and reach (as was flown at

a mission scale approximately 10x greater than Mars Surveyor program cost/launched mass), greater arm dexterity than that lenticular (extensible tube) 1-dof mechanism and actuated effector and more facility for on-arm instrumentation and *in situ* science. Our approach to meeting the volume constraint has been to develop a telescopic device (gas pressurization deployed) that efficiently collapses in stowage to about 1/5 of its two meter full extent. The particular workspace reach is dictated not only by typical science geometry, but also by need to access more pristine materials (i.e., far from an actively propelled lander's descent burn radius and chemical contaminants as possible). We have addressed reduction of mass through use of composite materials, the arm reported here having such a 8-segment deployable linkage, with a sampling effector currently under construction using similar materials and a new class of high power/mass density solid state piezoelectrically driven) motor. Note that the R&D prototypes we present here are terrestrial lab models -- devices and systems whose components assemblies are not space qualified to environment anti launch/flight dynamics. The graphite epoxy mechanical arm structures used have space heritage, and are extensible for flight model development. Laboratory actuators & drives are commercial 1) available brushed and magnetically encoded DC motors and harmonic/planetary gearing, all of which has space-qualified counterparts in various space platform deployment mechanisms, etc. Our primary goal in this R&D work has been to demonstrate new functionality and form-fit feasibility; we have addressed flight operations constraints & design requirements in a separate technology proposal to the Mars Surveyor '98 Announcement of Opportunity for landed science mission [4].

2.1 Arm Mechanical Design and Fabrication

The JPL arm, as shown in Figure 2 below in both stowed and deployed side-mounted lander configurations, is composed of two links, each link containing four telescoping tube segments. Each segment is graphite epoxy with a 30 cm. length and .040" wall thickness. The eight segments range from inner diameter of .525 to 1.75", in steps of .25". This allows the tube segments to nest inside of each other when stowed. Each segment has an internal and external fitting designed to lock the segments rigidly during deployment. A No. 1 Morse. taper lock, fabricated from Al 6061-T6, was used to provide the necessary locking force. Each complementary set of fittings was then bonded to the inside and outside of the tube segments using Stycast 9394 epoxy. When air pressure is applied, the outer fitting of one segment slides along the inner diameter of the next largest segment, eventually locking into its matching inner fitting. Each outer fitting has an embedded O-ring which seals each segment during locking. The ends of the first and last segment of each link have fittings (Al 6061-T6), which provide pressure seals for deployment, rubber grommets and clearance for cabling, and bolting surfaces for integration to the joint structure. Compressed air enters through the first fitting of Link 1 (upper arm) by means of a swage-lock right angle pressure fitting. The air is then transferred from Link 1 to Link 2 (forearm) by way of a "stinger" attachment. The "stinger" is attached to the top fitting of Link 2, and is sealed into Link 1 for deployment using a rubber grommet. The "stinger" detaches from Link 1 during actuation of the elbow joint after deployment. To prevent twisting of individual segments during deployment, aluminum alignment rails were bonded and screwed to the outer surface of the segments. The rail follows a broached groove in the internal fitting of the next largest segment, keeping the segments aligned until locking is complete. The arm links are integrated to the joint structures using English fasteners, and sealed using RTV Silicone.

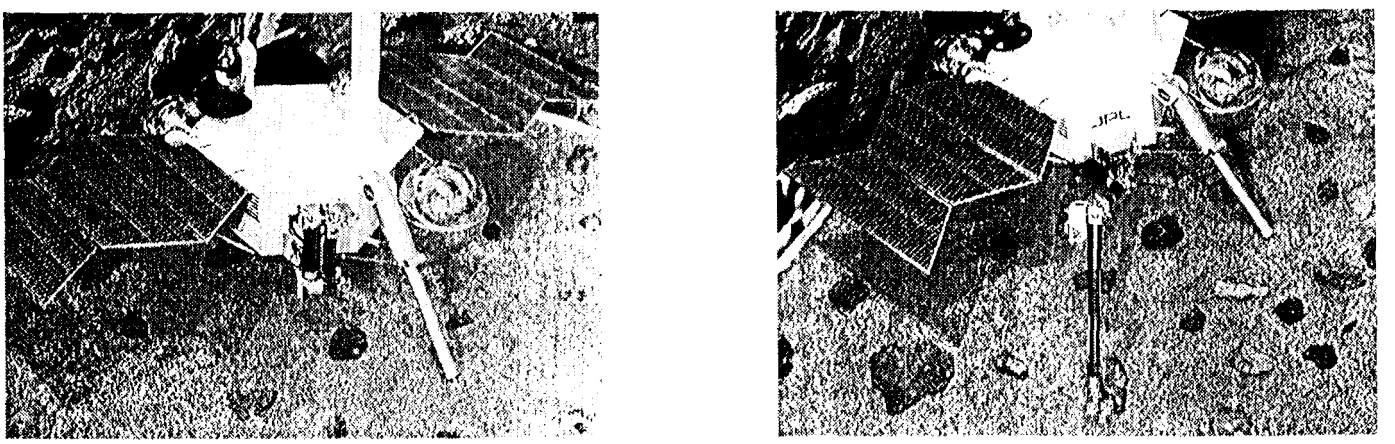

Figure 2 *JPL Lab Prototype Arm for Mars Lander Science Operations*
(left) *stowed arm configuration prior to gas driven deployment*
(right) *2.2 meter arm in a deployed operative configuration*

The joint mechanical components consisted of motors and harmonic drives at both the shoulder (DOF 1 and 2) and the elbow (DOF3). The motors used are Micro-Mo brushed DC motors with magnetic encoders and enclosed planetary gear reduction. The gear ratio at both shoulder locations is 134:1, while the ratio at the elbow is 43:1. These units are the inputs to harmonic drives that provide a further gear reduction ratio of 100:1, and have built in bearings. The motor/encoder/planetary units are cantilevered off of the harmonic drive by use of a fabricated A1 6061-T6 support.

The joint structural components consist of parts designed and fabricated for both the shoulder (DOF 1 and 2) and the elbow (DOF 3). The parts were fabricated by machining and welding from A1 6061-T6. The individual parts were assembled using English fasteners. The joints are then joined to their respective joint mechanical components by using metric fasteners to directly connect the fabricated hardware to both sides of the harmonic drive. The entire joint structure was then attached to the end fittings on Links 1 and 2 using English fasteners and RTV Silicone. When assembled, the manipulator arm structure is then connected to the Mars '98 lander model using four 1/4" bolts. A tie-down support was fabricated and attached to the lander mount to hold the free end of Link 2, which allows a symmetrical deployment of both arms simultaneously. The arm stand at 40 cm in height when stowed. For the demo, compressed air at 45 psi was used to deploy the arm. The total extension after deployment is approximately 2.2 meters.

2.2 Robot Arm Controls

We have implemented the arm servo and kinematic controls in a straightforward manner. The servo drive utilizes a set of Galil controller cards (4 dof/VME, 1 6-U slot, PI connector) addressing the Micro-Mo DC motors. Per above, the effective shoulder gearing is 13400:1, elbow 4300:1 -- respectively encoded to resolutions of (PPR: 10/channel, with quadrature, 40 counts per shaft rev) 11.7 uRadians/.000672 degs, and 6.5 uRadians/.00209 degs. Kinematics run in a VME-bus 68040/VxWorks environment. The 3-dof robot has a directly solvable inverse kinematic, and as sampling operations dictate, we can optionally control the arm in joint, simple coordinated joint, and full cartesian task space modes. Our emphasis on further controls development is two fold: 1) power optimal trajectory controls, reflecting realistic constraints of space operations (e.g., for a given task space goal, solve a trajectory for stated minimum power budget constraints over time [5]; 2) development of computationally simple but operationally robust behaviors [6, 7] for contact manipulation, task verification, and manipulator-health and safety (e.g., for a given set of operational sequences, and limited real-time sensor feedback, codify a useful set of finite state control actions). In general we are interested in those techniques and system architectures that facilitate bower- lution \cl)s~lr-let'~rct(c<[l autonomy and efficient supervisory interaction in telerobotic systems [8,9, and references therein].

3. ARM POSITION CONTROL RELATIVE TO TARGETS OF VISUAL INTEREST

With the above robot having been developed and placed under kinematic control, there remained a fundamental system problem -- establishing the relationship of the arm's position control inputs relative to a given world frame location -- e.g., as specified in multiple lander camera views by operator designation or machine vision. In its most familiar form, this problem is one of visual workspace calibration. Implemented in a traditional global way [10, 11], the solution requires explicit modeling of the camera parameters, robot arm kinematics, and estimation of non-homogeneous world frame transformations. This approach can be problematic if camera or arm information is lacking, or subject to unpredictable change (e.g., due to dimensional effects of extreme temperatures, operational system contingencies). In this section we describe an implementation approach yielding accurate estimates of an arm 3-D positioning control referenced directly to 2-D image plane information of multiple cameras. This *Camera Space Manipulation* method, as pioneered by Skaar and colleagues [12], leads to a highly efficient, robust means of localizing the arm relative to a task space coordinate (in generality, a position and orientation -- in the case of our 3-dof arm, an affixed world frame position -- e.g., a small rock sample in the reachable lander-field).

3.1 Technical Concept

The camera space manipulation(CSM) method [12] is predicated on estimating view parameters that describe the relationship between the appearance of image-plane visual features (or "cues") located on the manipulator and the internal joint configuration of the robot. Once this relationship is established, control of the robot is carried out entirely within the reference frame of the participating cameras without regard to any physical reference frame.

Using an orthographic camera model and a nominal model of the robot forward kinematic, one can relate the appearance of a cue on the manipulator end-effector to the internal joint configuration of the robot. Such cues are illustrated overleaf in **Figure 3**.

This relationship contains six view parameters which are to be determined. For each participating camera, these six view parameters are estimated initially by moving the robot through a "preplanned trajectory" which consists of a large set of joint configurations. At each joint configuration, the location of the cues on the end-effector in each of the participating camera's image planes (or camera space) is measured. Based on this large set of simultaneous joint and image plane samples, the six view parameters for each camera are estimated via a *nonlinear least-squared-error technique*. Implementation details of how to prescribe the preplanned trajectory are given in the section 3.3 below.

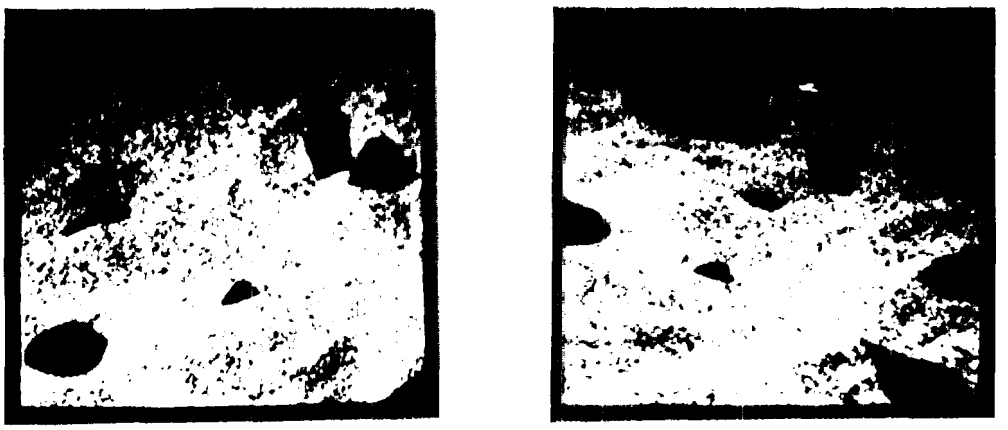

Figure 3 ***Operator Displays of Manipulator Cue for CSM Method***
(left) *view from boom mounted camera on top of lander*
(right) *view from camera mounted beneath solar panel*

For 3-D point positioning tasks using the three degree of freedom arm, the target point in physical space to be approached by the end-effector must be identified in each image plane of the participating cameras. As alluded above, for more general 3-D tasks involving six degree-of-freedom manipulators (where both position and orientation of the end-effector relative to the target object is required), multiple target points on the object to be approached must be identified.

Once the common target location has been identified *in* each camera's image plane and once the initial view estimates have been determined from the preplanned trajectory, the internal joint configuration that brings a point on the end-effector to the desired target point is determined by minimizing the sum of the squared difference between the image plane location of the target point and the image plane location of the point on the end-effector in both the x and y directions for each participating camera. Note that for general 3-D tasks, a minimum of two cameras must be used to resolve the joint configuration. Also, for a single target point, only three degrees-of-freedom can be determined using this method. Additional target points or other additional information must be used to determine the joint configuration for systems with more than three degrees-of-freedom.

Using the above approach, one can determine the joint configuration that takes the end-effector directly to the target point. However, due to various measurement and modeling errors present in the system, the final positioning precision achieved by the method may be not be acceptable. Instead, the robot is moved through a series of intermediate maneuvers that take the end-effector from a location far from the target position to the location which terminates at the target position. During these intermediate maneuvers, additional samples of the cue on the end-effector are acquired and the six view parameters are re-estimated using this new information. As the maneuver approaches the terminal target position, the weight placed on these visual samples is increased, thereby localizing the relationship described by the view parameters in both the camera space and joint space in which the maneuver terminates. This approach also helps to reduce the chance that the minimization routine used to resolve the joint configuration converges to the incorrect set of joint rotations. *This approach is one of the main reasons why the method of camera space manipulation is so robust to modeling and measurement errors*. Note the CSM method is not highly sensitive to obtaining visual information at each juncture during the approach to the target point. If cues are detected, then they are included in the estimates of the view parameters, otherwise the view parameters are not adjusted. The method is also very efficient with most computation intensive steps reserved to the image analysis routine. The cameras used for the system do not have to be very expensive or precise. Also, the system developed is very flexible in its setup.

There are certain conceptual similarities of CSM with autonomous guidance of robot arms by visual servoing [13]. Such methods control the manipulator based on determining the error between the current and desired locations of features on the manipulator in a closed-loop feedback manner. These methods typically require that all features to remain present in the image plane of the cameras and that the features be identified quickly so that the controller update to the robot can stay ahead of the motion of the robot. To our knowledge, relatively little has been published on experimentation in visual servoing for fully 3-D tasks involving both position and orientation of the robot's end-effector relative to a desired object.

3.2 Experimental Implementation

The experimental system, as configured for the general lander-based robotic arm scenario of Figure 1 and subsequent photo illustrations below, utilized he existing 3-dof arm as well as two low-cost CCD cameras mounted on the lander. One camera was positioned below one of the lander's solar panels and the other camera was mounted on a boom approximately 1 meter above the top of the lander base. The two cameras were separated by a distance of approximately 1.5 meters and both cameras were directed so that each camera viewed roughly the same region of the robot's workspace.

Computer hardware consisted of a 80386-based personal computer in which a Data Translation video framegrabber (DT2851) and a video multiplexer (used to allow more than one video input) were installed. Communication with the robot controller was carried out via the serial port on the computer. The camera space manipulation routines are all written in FORTRAN. Access to the framegrabber is also accomplished through FORTRAN subroutine calls to a library of functions provided by Data Translation. Additional code used to communicate with the robot controller through the serial port was written in Pascal.

A single white-centered cue was placed on the end-effector of the robot arm. This circular cue measured 1.5 inches in diameter and is constructed such that a dark ring (outside radius = 1.5 inches, inside radius = .75 inches) surrounds a white circle of radius .25 inches. The detection of this type of cue is accomplished by scanning a raw greyscale image acquired by the framegrabber for the black-white-black pattern associated with the white-centered cue. Note that the ratio of distances between the black and white regions of the cue will remain an invariant no matter the orientation of the camera viewing the cue. The circular cue does indeed map into the image as an ellipse, but this ratio of distances remains a constant. To detect a cue, the image is first searched horizontally for regions of the image that may contain a cue (i.e., constant ratio of distances). Once a possible cue is detected, this region of the image is searched again in the vertical direction and then in two diagonal directions. If the ratio of distances remains constant to within a certain tolerance level for each of these tests (horizontal, vertical, and two diagonal), then the possible cue is determined to be an actual cue. [Note: some difficulties with the cue detection routine described arose due to ambient-induced "false positives" from the simulated Martian surface regolith. To reduce such false positives" we tightened up on the tolerances used to judge if a detected cue was an actual cue. There of course exist more robust discriminable geometric cues that could be utilized at expense of modeling and computation).

3.3 Experimental Results

Each time the system was to be extensively tested, we performed a preplanned trajectory once at the beginning of the series of tests. This preplanned trajectory consisted of a set of 15 joint configurations which are specified *a priori*. This set of joint configurations was chosen such that the usable region of joint space that was to be used for the various tasks was, in some sense, spanned completely. Likewise, at these preplanned joint rotations, the cue on the end-effector of the robot also, in some sense, spanned completely the entire region of the image plane (or camera space) of the two cameras mounted on the lander. This requirement for spanning completely the joint region and image plane regions is necessary to estimating the six view parameters that describe the relationship between the appearance of the cues in each camera space and the joint configuration of the robot. If the preplanned trajectory is not chosen correctly, there is a danger in converging to the wrong set of view parameters, i.e. ones that only describe a small region of joint and camera space. This concept of completely spanning the range of input and output variables is a familiar concept in terms of parameter estimation theory. In general, the validity of the view parameters to completely describe the desired relationship can be evaluated by computing the residuals between the actual and predicted locations of the cues found during the preplanned trajectory. If these residuals are "reasonable" over the entire preplanned trajectory, then the view parameters are assumed to be correct. For this experimental setup, reasonable residuals were typically less than 15 pixels. While this is seemingly a large variance, note the ability of the orthographic camera model to accurately model the system when the cameras are relatively close to the robot arm is generally reduced.

Once the preplanned trajectory is completed, the user-prescribed target locations are then specified. This was accomplished by using the cursor routine (provided by the framegrabber) to choose a point on the desired target object in the image plane of each of the two cameras. In the past, cues were used to identify target locations in order to guarantee that a common point in physical space was detected in each camera. The fact that we employed *user-prescribed target locations in each image plane which may not correspond to a common physical point in physical space* was one of the novelties of the system developed.

After the target locations are prescribed in each camera the approach trajectory can be completed. This trajectory is determined by resolving the joint rotations as described in section 3.1. Instead of solving for the joint rotations that takes the end-effector directly to the target object, we resolve the joint rotations such that a series of three intermediate positions are achieved by the manipulator. At each intermediate position, the cue on the end effector is detected in each camera and the view parameters for each camera are re-estimated. This approach aids in further refining the view parameters such that they do a good job in capturing the relationship between the camera space and joint space in a region of camera space and joint space that is "local" to the final termination of the maneuver.

The first series of tests conducted with the system were used to *evaluate the actual positioning precision of the entire system.* During these tests, the actual physical target location was marked by a cue so that the inaccuracies associated with a non-compatible target point in each camera's image plane was not a factor. After conducting numerous tests in which the target location was moved through the region of the robot's workspace (as viewable from the cameras), the final positioning of a point on the end-effector was determined to be within 5 mm of the actual desired target location in all three physical space directions.

The next series of tests of the system dealt with the task of *inspecting various rock samples* which were stuttered around the lander. During these tests, the user of the system prescribed the target locations in each camera's image plane. Despite the fact that these target locations did not describe a common point in physical space, the system did not have any trouble converging to a set of joint rotations that would take the manipulator through the approach trajectory and then position the end-effector over the desired target object. (Evaluating the sensitivity of final positioning precision as a function of the incompatibility in the user-prescribed target locations is a substantive problem for further investigation.)

The final series of tests dealt with the *acquisition of a small rock sample* from the soil surface through the actuation of the powered scoop. For these tests, a point on the end-effector (or extension of the end-effector) was guided such that this point terminated over top of the sample to be acquired. Once this was accomplished the scoop was actuated so that the sample along with the soil itself was picked up by the scoop. These tests were also successful as illustrated in **Figure 4**.

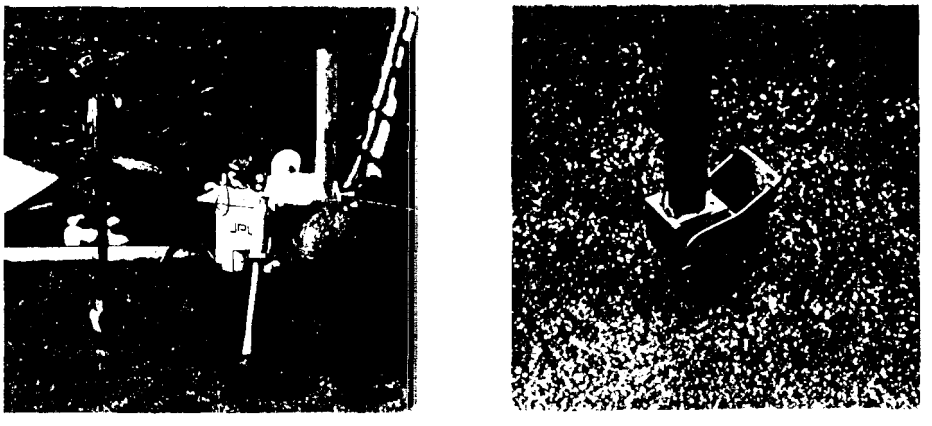

Figure 4 *Acquisition of a Small Rock Sample under CSM-based Position Control*
(left) *manipulator during approach to sample*
(right) *sample in actuated scoop just following terminal positioning & acquisition*

We note a few observant about the speed of the CSM operation. In the series experimental trials, the time the system took to identify the cue in both cameras and to resolve the joint rotations that take the robot to the next location in the approach trajectory was approximately 10 seconds. The image analysis took the majority of the computational effort. However, given the objective of

finding a single cue, we could localize cue search to a small region of each the camera's image plane. Using knowledge of the view parameters for each camera, one could predict location of the cue in the image plane, subsequently to center a small window about this predicted cue location and just search this window instead of the entire image plane. Also, the problem of false positives earlier noted is not very critical once the view parameters have been estimated. This is because a threshold on the accepted residual between the predicted and actual cue locations could be used to throw out any identified cues that did not conform to the current estimated view parameters. Typically this threshold was set at 5 pixels (i.e cues that had residuals larger than 15 pixels were not used to re-estimate the view parameters). As more and more approach trajectories were attempted and a large number of cues were detected, then typical residuals were on the order of 2-5 pixels. Certainly the entire computational process could be scaled as desired using a faster processor. Nor does the system need to stop in order to acquire a visual sample. If a video sample can be simultaneously acquired with samples of the joint rotations (via the motor encoders), then the method can proceed non-stop during its approach to the target position. With a quick processor, this approach may allow for many more cues to be detected during the trajectory that brings the end-effector to the desired target location.

3.4 Issues for Further investigation

In closing this section we comment briefly on some areas we believe useful to investigate further, as either might improve system robustness or generality of the techniques developed.

- *automatic correlation of a corresponding target location in one camera after selection in another*: Such an algorithm would circumvent the need for a user to prescribe matching target locations in both cameras as is done now. While the method of camera space manipulation does not rely on calibrated cameras, there is some information contained in the estimated view parameters and there should be a way to (recursively) use this information, along with possibly a very crude camera calibration to determine matching image plane locations that correspond to the same physical point.

- *camera space manipulation for obstacle/collision avoidance*: Currently the system possesses no capability to react to its environment -- i.e. the method only computes joint configurations that achieve the final positioning relative to the target object and assumes that the path that takes the robot there is obstacle free. For the task at hand, we envision the following type of scenario. In some manner, obstacles that are currently within the robot's workspace as viewed from the two cameras are identified in both image planes. This could be done by interactively outlining regions of camera space in which obstacles appear or by using computerized image analysis to estimate object edges and extremal contours. Then, in an off-line fashion, the joint rotations that bring the end-effector within the region of the obstacles would be computed. Finally, either as part of the CSM method or as something separate from CSM, the joint rotations that are computed during each trajectory would be modified using the information gained during the identification of the obstacles. In this manner, the trajectory computed during the approach phase could be constrained to occur within a certain region of the two image planes (or pinhole). The routine used to resolve the joint rotations based on the location of the target points and the current view parameter estimates putatively could include factors that are related to the identified obstacles. The final positioning would remain consistent, just the path followed being modified.

- *operator interaction and supervisory control*: [Two)111, It is possible to interactively, iteratively exploit most current knowledge of the view parameter estimates to show the operator how the manipulator will perform a certain task within the image planes of the two cameras. While the actual guidance of the manipulator may be slightly different due to the cues detected near the terminus of the maneuver, the simulation of the motion will be close to what is actually carried out. In this way, the operator can check the motion and make judgements about things like obstacle avoidance, reachability, etc.

4. ANALYSIS OF SAMPLES

The broader goal of work described in this paper is to couple efficient robotic activity to Mars lander-based science data acquisition and *in situ* analysis. There are many possible scenarios and enabling instruments for this pursuit [2]. We illustrate in **Figure 5** overleaf some of the robotic science options we have begun to evaluate in known correspondence with interests of the Mars science community. Certainly one area of major interest are those approaches which both spectrally and spatially discriminate information about the Mars surface constituents. Spectral information provide important clues to mineralogic classifications, and once samples of interest are so detected, then a microscopic visual analysis of sample surface structure (texturing, grains, local spectral properties, weathering) can reveal a great deal more of geologic, mineralogic and climatic importance, including assessments of volatile content with possible implications to exobiology and prior life on Mars. In this closing section, we briefly describe an area of develop-

ment being pursued by University of Arizona in collaboration with JPL colleagues: use of a multispectral handlens for spectroscopic analysis of the lander global near-field, and more detailed microscopic analysis of robotically acquired samples. The lander-based operational configuration for this work is illustrated in **Figure 6** overleaf.

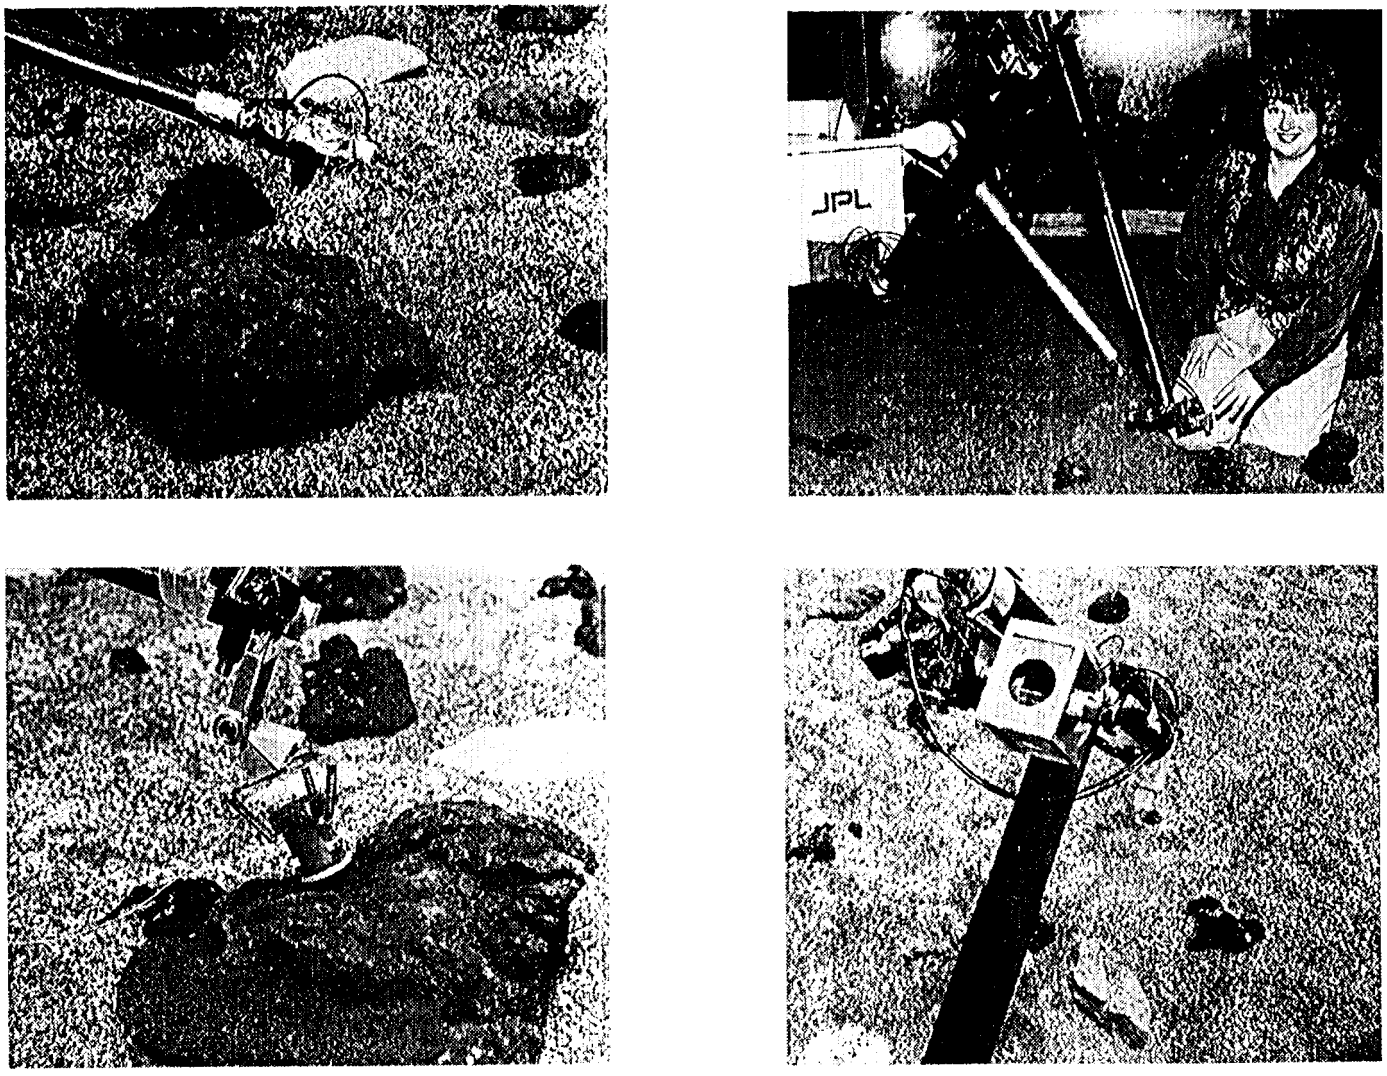

Figure 5 ***Simulated Mars Lander Science Operations***

(top left) *microscopic viewing of large rock sample via end-arm mounted camera*

(bottom left) *emplacement of simulated alpha proton x-ray device to rock sample*

(upper right) *end arm mounted point spectrometer for local sample analysis (D. Blaney)*

(bottom right) *opposed jaw dual actuated effector in terminal approach to sample acquisition*

The Multispectral Handlens (MSHL) is being developed at the Laboratory for Autonomous Remote Sensing at the University of Arizona as part of an integrated smart system for robotic sampling and analysis of Mars [14]. The MSHL functions as a lander-based spectral imager to obtain spatial and compositional data in the far-field, near-field (local lander environment), and microscopic regimes. Spectroscopic machine vision provides a wealth of compositional information compared with monochrome or 3-color systems. Such an instrument is necessary to perform geologic science-based targeting and verification for arm manipulation and sampling. Our goal is to develop a fully integrated intelligent system that will have the capability to image a scene, autono-

mously perform initial mineralogic classification based on the spectral and spatial environment, determine which elements merit closer examination, and collect samples via the robotic arm for further analysis using a lander based suite of instruments.

Our approach is (0 draw on multiple disciplines to develop a suite of practical capabilities for autonomous non-contact optical compositional determination in a spatial context (i.e., determining "what is it and where is it".)" Interpretation of such complicated and often voluminous data is presently a time-consuming process requiring highly specialized human analysts. We believe that the information returned from an autonomous sensor system should be a targeted, context-sensitive, high-level extract of the more voluminous data obtained by the system. This is true whether the recipient is another machine or is human, especially if real-time results are required. The sensor system must be able to adaptively decide what data to take, how and when to take it, and how to process and analyze it to suitably return the desired information. Such a system requires automated intelligence, on-board processing and analysis, and adaptive decision-making capability. These objectives are particularly challenging in a highly unstructured environment such as a natural geologic setting.

For the role of spatially and compositionally mapping the local environment, the MSHL will be mounted on a pan-tilt head on a fixed mast above the simulated lander. It can also be mounted as a fixed lander-based instrument for use as a multispectral microscope. In the future, improved lighter and smaller versions of the current instrument could be placed directly on the end effecter of the robotic arm, or on a moveable mast for use both in near field and microscopic regimes. The current prototype MSHL provides spectral imaging from 0.4-1.0 mm wavelength using 9 interference filters, over distances from infinity to a few centimeters. At long ranges it can characterize the general geology of the landing site. At closer ranges, it can map the local environment spectrally and compositionally for use in targeting and verification for arm manipulation and sampling. At very close range (a few cm) the instrument functions as a multispectral microscope, returning spectral and spatial information for individual mineral grains down to ~20 microns/pixel. We have had good success at spectral extraction and mineralogic identification over all of these regimes, an illustrative example being given by Figure 6. We have begun integration with the lander robotic arm using the MSHL as a representative automated scientific instrument. Samples acquired by the robotic arm were deposited on a lander-mounted turntable which then rotated the samples under the MSHL. Such a delivery system can be extended to transport samples to several instruments without requiring further arm movement.

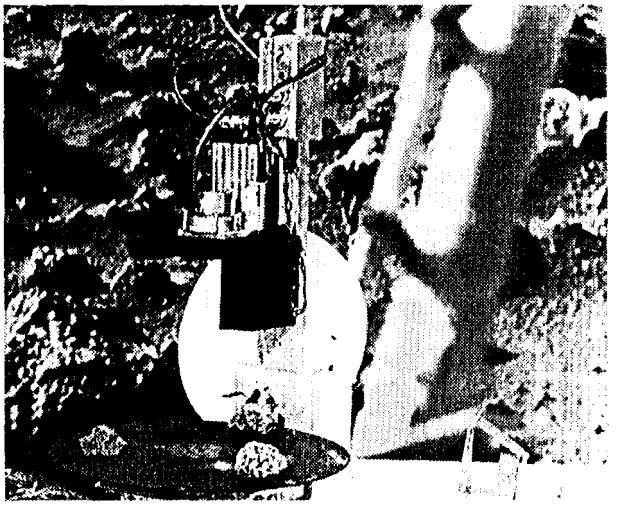

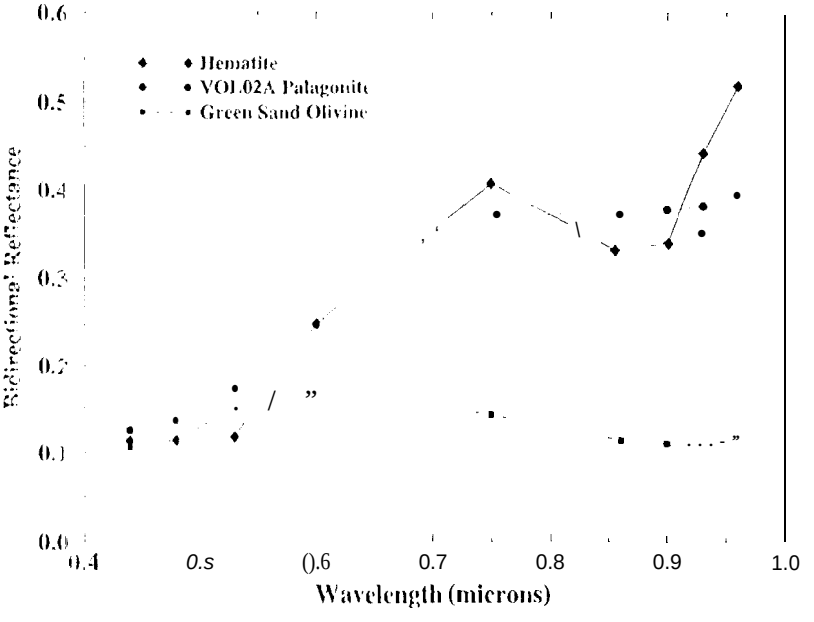

Figure 6 *Multispectral Handlens and Lander Based Operations*

(left) University of Arizona Multispectral Handlens (R. Singer et al.) in deck mount configuration

(right) spectra extracted from minerals believed to exist on the Martian surface. Subtle spectral differences in the visible (0.4-0.7 microns) are complemented by characteristic spectral features in the near-IR (0.7-1.0 microns) providing the capability for compositional discrimination and mineralogic identification. This data highlights the necessity of multispectral imagery over that of monochrome or 3-color systems

An important feature of the MSHL is the ability to perform self-contained autofocusing. The autofocusing algorithm analyzes and optimizes the information content of the image, eliminating the need for active ranging systems. Autofocus is necessary since the camera is designed to function from infinity down to the microscopic regime. A fixed-focus system would be inadequate over this range. Autofocus also eliminates the need for precise positioning of irregular geologic samples by either a robotic arm or a lander-mounted sample delivery system.

5. ACKNOWLEDGEMENTS

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The authors acknowledge use of the opposed jaw dual actuated effector used in some of the reported experimentation (per Figure 5 bottom right as developed by A. Ohm and R. Volpe, JPL Rover Technology task).

REFERENCES

1. R. E. Arvidson, J. L. Gooding, and H. J. Moore, "The Martian Surface as imaged, sampled, and analyzed by the Viking landers," Rev. Geophys. Space Phys., Vol 27, pp. 39-60, 1989.

2. *Mars Surveyor Science Objectives and Measurements Requirements Workshop* (Eds. D. J. McCleese, S. W. Squyres, S. E. Smrekar, and J. B. Plescia), JPL Tech. Report No. D12017, Jet Propulsion Laboratory, Pasadena, CA. (May 10-12, 1994).

3. R. B. Seger and V. P. Gillespie, "The Viking surface sampler," in Proc. 8th Aerospace Mechanisms Symposium, pp. 235-245, Oct. 18-19, 1973.

4. NASA AO 95-0 SS-03, blurs Surveyor Program '98 Lander and Orbiter Mission Missions Science Instruments; see also related Proposal Information Packages (Vol 1: Lander Science Definition Team Report, Vol 2: 1998 Lander Mission), open to August 8, 1995.

5. S. Dubowsky, C. Moore, and C. Sunuda, "On the design and task planning of power efficient field robotic systems," in Proc. ANS 6th Topical Meeting on Robotics and Remote Systems, pp. 501 - 508, Feb 1995, Monterey, CA.

6. R. A. Brooks, "Intelligence without representation," in *The Foundations of Artificial Intelligence* (Ed., D. Kirsch). The MIT Press, 1992: Cambridge, MA, and further references therein, D. McFarland and T. Bosser, *Intelligent Behavior in Animals and Robots*. The MIT Press, 1993: Cambridge, MA.

7. M.R. Stein and R.P. Paul, "Operator interface for time-delayed teleoperation, with a behavior-based controller," in Proc. IEEE Intl. Conf. Robotics & Automation, San Diego, CA, May, 1994; Matthew R. Stein, Behavior-Based Control for Time-Delayed Teleoperation, Ph.D. Thesis (April, 1994), The University of Pennsylvania, Philadelphia (Advisor: R.P. Paul; NASA Technical Advisor: P. S. Schenker, Jet Propulsion Laboratory); M. Stein, R. Paul, E. Paljug, and P. Schenker, "A cross-country teleprogramming experiment," in Proc. 1995 IEEE-RSJ Intl. Conf. IROS, Pittsburgh, PA, August.

8. P. S. Schenker, "Intelligent robots for space applications," pp. 545-591 in *Intelligent Robotic Systems: Analysis, Design, and Programming* (S. Tzafestas, Ed.). Marcel Dekker, 1991: New York, NY; P. S. Schenker, "NASA research & development for space robotics," IEEE Trans. Aerospace Electr. Sys (special section on space telerobotics, R. Volz, Ed.), Vol. 24, no. 5, pp. 523-534, September 1988.

9. P. S. Schenker and G. T. McKee, "Man-machine interaction in telerobotic systems and issues in the architecture of intelligent and cooperative control," in Proc. IEEE/Intl. Symp. Intell. Control Workshop (Architectures for Semiotic Modeling and Situation Analysis in Large Complex Systems, Orgs.: J. Albus, A. Meystel, D. Pospelov, T. Reader). Monterey, CA, August, 1995.

10. B. K. P. Horn, *Robot Vision*, The MIT Press (McGraw-Hill Cambridge, MA, 1986.

11. R. Y. Tsai, "A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses," IEEE Jrnl. Robotics and Automation, Vol. RA-3, No. 4, pp. 323-344, August 1987.

12. S.B. Skaar, W.H. Brockman, R. Hanson, "Camera-Space Manipulation," International Journal of Robotic Research, Vol. 6, No. 4, pp. 20-37, 1987; S.B. Skaar, W.H. Brockman, W.S. Jang, "Three dimensional Camera Space Manipulation," international Journal of Robotic Research, Vol. 9, No. 4, pp. 22-39, 1991; S.B. Skaar, I.F.E. Gonzalez-Galvan, "Versatile and precise vision-based manipulation," *Teleoperation and Robotics in Space*, S.B. Skaar and C.F. Ruoff, eds., Chapter 9, AIAA Press, Washington, September, 1994.

13. J.T. Feddema, C.S. George Lee, O.R. Mitchell, "Weighted selection of image features for resolved rate visual feedback control," IEEE Trans. on Robotics and Automation, Vol. 7, No. 1, pp. 31-47, Feb. 1991; L.E. Weiss, A.C. Sanderson, C.P. Neuman, "Dynamic sensor-based control of robots with visual feedback," IEEE Trans. on Robotics and Automation, Vol. 3, No. 5, pp. 404-417, Oct. 1987; B. Espiau, F. Chaumette, P. Rives, "A new approach to visual servoing in robotics," IEEE Trans. on Robotics and Automation, Vol. 8, No. 3, pp. 313-326, June 1992; *Visual Servoing: Achievements, Applications, and Open Problems* (Orgs., G. Hager and S. Hutchinson), Workshop Proc. of 1994 IEEE Intl. Conf. Robotics and Autom. (M-5), San Diego, CA, May.

14. R. J. Reid and R. B. Singer, "Mineral identification from well-calibrated multispectral imaging on the surface of Mars," Lunar and Planet. Sci. Conf. 26, pp. 1155-1156. (abstract), 1995; also R.B. Singer, P. Lever, and S. Larson, "Autonomous near-field remote sensing, for planetary exploration, 12 pp. unpublished proposal), 1994.